[To be read before The Royal Statistical Society at a meeting organised by the Discussion Meetings Committee to be held at 12 Errol Street, London EC1Y 8LX on Thursday, 16 June 2022, with Dr Shirley Coleman in the Chair.]

Semi-Mechanistic Bayesian modeling of COVID-19 with Renewal Processes

Samir Bhatt[1], Neil Ferguson[1], Seth Flaxman[2], Axel Gandy[2,†], Swapnil Mishra[1], and James A. Scott[2]

[1]MRC Centre for Global Infectious Disease Analysis, Imperial College, London, UK
[2]Department of Mathematics, Imperial College, London, UK
† Corresponding author: a.gandy@imperial.ac.uk

Abstract

We propose a general Bayesian approach to modeling epidemics such as COVID-19. The approach grew out of specific analyses conducted during the pandemic, in particular an analysis concerning the effects of non-pharmaceutical interventions (NPIs) in reducing COVID-19 transmission in 11 European countries (Flaxman et al., 2020b). The model parameterizes the time varying reproduction number R_t through a multilevel regression framework in which covariates can be governmental interventions, changes in mobility patterns, or other behavioural measures. Bayesian multilevel modelling allows a joint fit across regions, with partial pooling to share strength. This innovation was critical to our timely estimates of the impact of lockdown and other NPIs in the European epidemics: estimates from countries at later stages in their epidemics informed those of countries at earlier stages. Originally released as Imperial College Report 13 Flaxman et al. (2020a) on 30 March 2020, the validity of this approach was borne out by the subsequent course of the epidemic. Our framework provides a fully generative model for latent infections and derived observations, including deaths, cases, hospitalizations, ICU admissions and seroprevalence surveys. One issue surrounding our model's use during the COVID-19 pandemic is the confounded nature of NPIs and mobility. We explore this issue using our R package `epidemia` which implements the approach in Stan. Versions of our model were used in an ongoing way by New York State, Tennessee and Scotland to estimate the current epidemic situation and make policy decisions.

1 Introduction

This article presents a general framework for semi-mechanistic Bayesian modeling of infectious diseases using renewal processes. The term semi-mechanistic relates to statistical estimation within some constrained mechanism. Variants of this general model have been used in specific analyses of Covid-19 (Flaxman et al., 2020b; Vollmer et al., 2020; Mellan et al., 2020; Unwin et al., 2020a; NYS Press Office, 2020; Olney et al., 2020; The Scottish Government, 2020; Mishra et al., 2020), and continue to be used in ongoing work to make policy decisions. The present article motivates and discusses the key statistical and epidemiological features of this framework, starting from a counting process setup. Various extensions of the basic model are considered, including a latent infection process. We discuss limitations and applications of the modelling framework to stimulate further research.

The model uses a flexible regression-based framework for parameterising transmission and ascertainment rates. This allows the fitting of multilevel models (Gelman and Hill, 2006; Hox et al., 2010; Kreft and de Leeuw, 2011) for several regions simultaneously. Such partial pooling of parameters has specific advantages in the context of infectious diseases. Suppose we wish to estimate the effect of NPIs (Cowling et al., 2020; Flaxman et al., 2020b; Islam et al., 2020) or mobility (Badr et al., 2020; Miller et al., 2020) on transmission rates. Estimating separate models for each region could lead to a set of noisy estimates for at least two reasons. There is typically little high quality data at the early stages of an epidemic, and such data is generally correlated, reducing the information content that can be used to infer such an effect. In addition, NPIs often occur in quick succession and their effects are confounded (of HKSAR, 2003; WHO,

2003). This is exacerbated by the random times between infections (the generation distribution) and between infections and observations, which smooths the observed data, making it more difficult to attribute changes in transmission rates to a particular NPI. Alternatively, one could fit a single model, pooling the effect across all regions. This ignores region-level variation and can lead to poor predictive performance, in particular underestimating variance for previously unmodeled regions. One could augment such a model with group-level indicators, but this results in a large number of parameters, which are difficult to estimate and leads to overfitting with classical estimation techniques. The Bayesian multilevel regression approach of partial pooling provides a natural solution to these problems.

Sometimes the inferential goal is not to assess the effect of a covariate on outcomes, but rather to infer *latent* transmission rates and their effect on outcomes. Previous studies have focused on estimating reproduction numbers from case data (Ferguson et al., 2001; Riley et al., 2003; Bettencourt and Ribeiro, 2008; Fraser et al., 2009; Kelly et al., 2010; Cori et al., 2013), sometimes directly substituting observed case counts for the unknown number of infected individuals (Wallinga and Teunis, 2004). However, the emergence of SARS-CoV-2 has highlighted shortcomings of methods that rely on just case data. Limited testing capacity at the early stages of the pandemic led to only a small proportion of infections being detected and reported (Li et al., 2020). Those tested were typically more likely to have been hospitalised or were at higher risk of infection or death. In particular this proportion, referred to as the infection ascertainment rate (IAR) is country-specific and likely to have changed over time due to changes in testing policies and capacity. If unaccounted for, it will lead to biases in the inferred transmission rates.

The problem of varying case ascertainment highlights the need for more flexible observation models to rigorously incorporate various types of data, from hospitalizations to seroprevalence surveys. Daily death data has been used in (Flaxman et al., 2020b) to recover reproduction numbers in the early stages of the SARS-CoV-2 pandemic, and has been seen as more reliable than case data. However, there have been clear variations in case and mortality definitions as well as reporting across time and countries. It is therefore important to appropriately model noise within the the observational models. Our framework allows for multiple types of data including deaths, cases, hospitalizations, ICU admissions and the results of seroprevalence surveys. This improves robustness of inferred parameters to biases in any one type of data.

Our model uses discrete renewal processes to propagate infections within modeled populations. These have been used in a number of previous studies (Fraser, 2007; Cori et al., 2013; Nouvellet et al., 2018; Cauchemez et al., 2008), and are linked to other popular approaches to infectious disease modeling. (Champredon et al., 2018) show that the renewal equation leads to identical dynamics as Erlang-Distributed Susceptible-Exposed-Infected-Recovered (SEIR) compartmental models, when a particular form is used for the generation distribution. A special case of this is the standard Susceptible-Infected-Recovered (SIR) model (Kermack, William Ogilvy and McKendrick, 1927). The approach is also connected to counting processes such as the general branching processes(Bellman and Harris, 1948; Kimmel, 1983; Jagers, 1969; Crump and Mode, 1969; Pakkanen et al., 2022). Self exciting Hawkes processes are also related to renewal processes, with the expectation of the Hawkes intensity function resulting in a renewal equation (Rizoiu et al., 2017).

We describe the general model in detail, and start by considering the simplest version in Section 2. The motivation for the model lies in continuous-time counting processes, and this connection is discussed in Section 3. Sections 4 and 5 present the infection and observation processes in more detail, and consider important extensions of the basic model. Section 6 considers how to use the framework for multilevel modeling. Section 7 compares our approach to standard time series models, and outlines the key challenges involved in modeling with our framework. Section 8 considers the specific aspect of confounding and causality when estimating the effects of variables on transmission rates. Section 8.4 considers a simple simulation analysis and a discussion of when our approach is expected to fail, and Section 9 has a brief discussion.

2 Model Overview

We now formulate a basic version of the model for one homogeneous population. The same model can be used for multiple regions or groups jointly. In what follows we will consider discrete time, $t \in \mathbb{Z}$, but continuous analogues are similar (Pakkanen et al., 2022). Let $R_t > 0$ be the time-varying reproduction number at time $t > 0$, determining the average number of secondary infections caused by an infected person. The number of seeded infections $i_v, \ldots, i_{-1}, i_0$

for some integer-valued timepoint $v \leq 0$ are given a prior distribution. For $t > 0$, we let new infections i_t be defined by

$$i_t = R_t \sum_{s<t} i_s g_{t-s}, \tag{1}$$

where the generation time, the lag between infections, is given through a probability mass function g, i.e. $g_t \geq 0$ and $\sum_{t=1}^{\infty} g_t = 1$.

We never observe the exact time at which a person becomes infected; rather we observe events reported by health care systems, introducing bias due to under/over ascertainment and reporting lags. Recorded observations occur at certain times $t > 0$. In general, there may be multiple types; cases, hospitalisations and death counts for example. Each such type is driven by its own time-varying ascertainment rate $\alpha_t > 0$. The predictions of the recorded observation at time t is linked to past infections by

$$\hat{y}_t = \alpha_t \sum_{s \leq t} i_s \pi_{t-s}, \tag{2}$$

where π is a distribution for the lag between an infection and when it gives rise to a recorded observation. The sampling distribution of the observations with these means is typically nonegative and discrete, and may depend on auxiliary parameters. When multiple types are observed, we can superscript the quantities as $\hat{y}_t^{(l)}, \alpha_t^{(l)}$ and $\pi^{(l)}$ and assign independent sampling distributions for each type. We can connect $\hat{y}$ to a likelihood for true observations y with, e.g. Poisson or Negative Binomial likelihoods.

Transmission rates R_t and ascertainment rates α_t can be modeled flexibly using Bayesian regression models. Multilevel modeling allows us to share parameters between regions, borrowing strength from regions with advanced epidemics to inform estimates in regions with earlier epidemics. One can, for example, model transmission rates as depending on a binary covariate for an NPI, say full lockdown. The coefficient for this can be *partially pooled* between these groups. The effect is to share information between groups, while still permitting between group variation.

3 Motivation from continuous time

Our model can be motivated from a continuous time perspective as follows. Infections give rise to additional infections in the future, referred to as offspring. Letting $N^I(t)$ denote the number of infections occurring up to time t, defined by its intensity

$$\lambda(t) = R(t) \int_{s<t} g(t-s) N^I(ds), \quad t > 0, \tag{3}$$

where g is the density of a probability distribution on $\mathbb{R}^+$ defining the time between infections, and where $\{R(t) : t > 0\}$ is a non-negative stochastic process. The process can be initialised by assuming values for $N^I(t)$ for t in the seeding period $[v, 0]$.

Equation (3) is similar to the Hawkes intensity, however the *memory kernel* g is scaled by a time-specific factor $R(t)$. The integrand g allows the intensity to increase due to previous infection events, while $R(t)$ tempers the intensity for other time-specific considerations. Under this assumption, since g integrates to unity, the expected number of offspring is simply $R(t)$, and so this is the *instantaneous reproduction number* or alternatively the *branching factor* of the Hawkes process. The generation time, defined as the time from an infection to a secondary infection, is distributed according to g and so g is the *generation distribution*.

Recorded observations are caused by infections that occurred in the past - that is, a given infection may lead to observation events (cases or deaths) in the future. Letting $N^Y(t)$ be the count of some observation type over time defined by the intensity

$$\lambda_y(t) = \alpha(t) \int_{s<t} \pi(t-s) N^I(ds), \tag{4}$$

for $t > 0$, where $\pi : \mathbb{R}^+ \to \mathbb{R}^+$ is a function and $\{\alpha(t) : t \geq 0\}$ is a non-negative stochastic process. This is similar to Equation (3), however the intensity increases due to past infections, rather than past observations.

Consider the special case where π is a probability density and where $\alpha(t') = \alpha(t)$ for all t'. The average number of observation events attributable to a single infection is then $\alpha(t)$, and so this is an *instantaneous ascertainment rate*. π is then interpreted as the distribution for the time from an infection to an observation, and therefore we call it the *infection-to-observation* distribution.

4 Infection Process

Starting from the continuous model, we now describe a discrete model, which results in the formulation of Section 2. This discrete model is more amenable to inference. Let I_t be the number of new infections at time t; this is the equivalent of $N_t^I - N_{t-1}^I$ in the continuous model. As a basic modelling block we use the following discrete version of (3):

$$\mathrm{E}[I_t | R_{1:t}, I_{v:t-1}] = R_t L_t, \tag{5}$$

where $L_t := \sum_{s<t} I_s g_{t-s}$ is the *case load* or *total infectiousness* by time $t > 0$. Moreover, letting $i_t := \mathrm{E}[I_t | R_{1:t}, I_{v:0}]$ and taking the conditional expectation given reproduction numbers $R_{1:t}$ and seeded infections $I_{v:0}$ on both sides of (5) gives

$$i_t = R_t \,\mathrm{E}[L_t | R_{1:t}, I_{v:0}] = R_t \sum_{s<t} \mathrm{E}[I_s | R_{1:s}, I_{v:0}] g_{t-s} = R_t \sum_{s<t} i_s g_{t-s},$$

which is Equation (1). This is a discrete renewal equation, which can alternatively be interpreted as an AR(t)-process with known coefficients g_k. From this point of view, the basic model in Section 2 uses i_t as synonymous with actual infections. Since infections are simply a deterministic function of other parameters, there is no need to treat them as unknown latent parameters to sample. This can lead to lower sampling times and faster convergence when performing Bayesian inference.

4.1 Modeling Latent Infections

The model of Section 2 can be extended by replacing each i_t with the actual infections from the counting process I_t, and then assigning a prior to I_t. Although sampling can be slower, this has certain advantages. When past infection counts are low, significant variance in the offspring distribution can imply that the number of new infections I_t has high variance. This is not explicitly accounted for in the basic model. In addition, this approach cleanly separates infections and observations; the latter being modeled *conditional* on actual infections. The sampling distribution can then focus on idiosyncrasies relating to the observation process.

We assign a prior to I_t conditional on previous infections and current transmission R_t. The expected value for this is given by Equation (5). Appendix 10.1 shows that assuming the variance of the prior to be a constant proportion d of this mean is equivalent to letting d be the *coefficient of dispersion* for the offspring distribution. $d > 1$ implies overdispersion, and can be used to account for super-spreading events, which has been shown to be an important aspect for modeling Covid-19 (Lloyd-Smith et al., 2005). The parameter d can be assigned a prior.

Any two parameter family can be used to match these first two moments. Letting this be continuous rather than discrete allows inference to proceed using Hamiltonian Monte Carlo, whereby new values for I_t are proposed simultaneously with all other parameters. Possible candidates include log-normal, gamma and the Weibull distributions. If an explicit distribution for the offspring distribution is desired, one can show that assuming a Gamma distribution with rate λ for this results in a Gamma distribution for I_t with rate λ. The coefficient of dispersion is then simply $d = \lambda^{-1}$.

4.2 Population Adjustments

If R_t remains above unity over time, infections grow exponentially without limit (in branching processes literature this is referred to as super-critical). In practice, infections should be bounded from above by S_0, the initial susceptible population. All else being equal, transmission rates are expected to fall as the susceptible population is depleted.

Consider first the model using I_t, which was described in Section 4.1. Equation 5 can be replaced with

$$\mathrm{E}[I_t|R_{1:t}, I_{v:t-1}] = (S_0 - I_{t-1})\left(1 - \exp\left(-\frac{R_{u,t}L_t}{S_0}\right)\right), \tag{6}$$

where $R_{u,t}$ is an *unadjusted* reproduction number, which does not account for the susceptible population. This satisfies intuitive properties. As the *unadjusted* expected infections $R_{u,t}L_t$ approaches infinity, the *adjusted* expected value approaches the remaining susceptible population. The motivation for and derivation of Equation (6) is provided in Appendix 10.2. In short, this is the solution to a continuous time model whose intensity is a simplification of Equation (3). We must also ensure that the distribution of I_t cannot put positive mass above $S_0 - I_{t-1}$. A simple solution is to use truncated distributions. Of course, this adjusts the mean value from Equation (6), however this is unlikely to be significant unless the susceptible population is close to zero.

In the basic model, one can apply the adjustment to i_t by replacing L_t in Equation (6) with

$$\mathrm{E}(L_t|R_{1:t}, I_{v:0}) = \sum_{s<t} i_s g_{t-s}. \tag{7}$$

5 Observations

Observations are modeled in discrete time, analogous to how we treated infections in Section 4. Letting $\pi : \mathbb{N} \to \mathbb{R}^+$ and $Y_t := N_t^Y - N_{t-1}^Y$, the discrete analogue to Equation (4) is

$$\mathrm{E}[Y_t|\alpha_t, I_{v:t}] = \alpha_t \sum_{s\leq t} I_s \pi_{t-s}. \tag{8}$$

Taking the expected value of the above given seeded infections, transmission rates and the current ascertainment rate gives

$$\mathrm{E}[Y_t|\alpha_t, R_{1:t}, I_{v:0}] = \alpha_t \sum_{s\leq t} i_s \pi_{t-s}, \tag{9}$$

which is recognisable as Equation (2). Thus we have two possible expressions for the mean of Y_t, one given actual infections, and the other given expected infections i_t. The basic model of Section 2 uses the latter, while the extension in Section 4.1 uses the former.

We assume that $Y_t \sim \mathcal{F}(y_t, \phi)$, where $\mathcal{F}$ is a non-negative discrete family parameterised by its mean y_t and potentially an auxiliary parameter ϕ. This could be a Poisson distribution, where there is no auxiliary parameter. Using a quasi-Poisson or negative binomial instead allows for overdispersion. This can be useful to capture, for example, day-to-day variation in ascertainment rates when infection counts are low. The mean y_t can be taken to be either (8) or (9), the latter being used in the basic version of the model. Hidden in this formulation is the assumption that the Y_t's are conditionally independent given y_t. Using multiple observation series $Y_t^{(l)}$ can help to improve the model inferences and identifiability of certain parameters. We simply assume that each such series is conditionally independent given the underlying infection process.

6 Multilevel Models

Transmission rates can be modeled quite generally within the framework. If the aim is simply to estimate transmission in a single region over time, one approach could be to let $R_t = \psi^{-1}(\gamma_t)$, where ψ is a link function and γ_t is some autocorrelation process, for example a random walk. If the goal is to estimate the effect of NPIs in M regions on transmission, we can let $R_t^{(m)}$ denote the time-varying reproduction number in region m at time t, specifying

$$R_t^{(m)} = \psi^{-1}\left(\beta_0^{(m)} + \sum_{l=1}^{p} x_t^{(m)} \beta_k^{(m)}\right), \tag{10}$$

where $x_t^{(m)}$ are binary encodings of NPIs, and $\beta_0^{(m)}$ and $\beta_k^{(m)}$ are region-specific intercepts and effects respectively. The intercepts are used to allow regions to have their own baseline transmission rates (and can be interpreted as R_0). Collecting these group specific parameters into $\beta^{(m)}$, we can partially pool them by letting $\beta^{(m)} \sim \mathcal{N}(0, \Sigma)$, for each group m, and then assigning a prior to the covariance matrix Σ. This could be an inverse-Wishart prior, or alternatively, Σ can be decomposed into variances and a correlation matrix, which are each given separate priors (Tokuda et al., 2011).

One possible option for ψ is the log-link. This provides easily interpretable effect sizes; a one unit change in a covariate multiplies transmission by a constant factor. However, this can lead to prior mass on unreasonably high transmission rates. With this in mind, an alternative is to use a generalisation of the logit link for which

$$\psi^{-1}(x) = \frac{K}{1 + e^{-x}}, \tag{11}$$

and where K is the maximum possible value for transmission rates. This serves a similar purpose to the carrying capacity in a logistic growth model.

The ascertainment rate α_t can also be modeled with similar considerations to the above. This flexibility is useful, particularly because these quantities are likely to change as an epidemic progresses. This has been clearly seen during the Covid-19 epidemic, where the infection ascertainment rate changed over time due to increased testing capacity and improved contact tracing systems. Multilevel modeling approaches are equal applicable to the specification of α_t.

7 Forecasting, epidemiological constants, and seeding

A key benefit of using a semi-mechanistic approach is that forecasts are constrained by plausible epidemiological mechanisms. For example, in the absence of any further interventions or behavioural changes, and looking at a medium term forecast of just incidence (daily new cases/infections), a traditional time series forecasting approach may predict a constant function based on observing broadly constant incidence, but a semi-mechanistic approach would expect a monotonic decrease based on a constant rate of transmission and the depletion of the susceptible population. The performance of epidemiologically constrained models is generally good (Carias et al., 2019); this is perhaps not surprising as examining the discrete renewal equation shows that these models correspond to autoregressive filters with a convex combination of coefficients specified by the generation distribution. However, similar to financial forecasting, the predictive capability of epidemic models are likely to be better interpreted as scenarios rather than actual predictions due rapidly changing policies and the unpredictable behavioural responses of human populations.

A second benefit of epidemic models is to provide a plausible mechanism to explain (non causally) the changes observed in noisy data. For example, in estimating the effect of an intervention on observed death data, we need to consider what that intervention affects, i.e. the time-varying reproduction number R_t. As we have described above, we link the reproduction number to the number of latent infections to an observed quantity (cases, hospitalisations, or deaths) with an epidemiologically motivated mechanism. While we can statistically estimate parameters for how the intervention affects R_t, certain important parameters will be entirely unidentifiable and need to be fixed as constants or with very tight priors. For example, to reliably estimate the number of infections, an infection fatality ratio needs to be chosen. A failure to choose an appropriate infection fatality ratio can result in a bimodal posterior where changes can either be attributed to rapid depletion of the susceptible population (so-called "herd immunity") or to the effect of interventions. From a statistical perspective, it is difficult to disentangle which mode of the posterior best represents reality. When properly interpreted, this can be an informative finding, but to obtain epidemiologically plausible estimates from the semi-mechanistic model requires fixing the infection fatality ratio using estimates obtained from the literature. A second example is the fact that Covid-19 deaths occur, on average 3 weeks after infection. Omitting the infection-to-symptom and symptom-to-death distributions will bias effect estimates. This point was proved over and over when rising cases at the beginning of an epidemic wave were dismissed by the claim that hospitalisations and deaths were not rising, forgetting about the inherent lag in these measures.

Infection seeding is another fundamentally challenging aspect of epidemic modelling. Estimating the initial effect of seeding is crucial to understanding a baseline reproduction number (R_0) which is modified by behaviour, interventions, and population depletion. This seeding is heavily confounded by importation and underascertainment. Both these factors

can influence estimates of the initial growth rates, and this in turn can affect the impact of changes in transmission as time progresses. We have proposed heuristic approaches to mitigate issues with early seeding, but principled statistical approaches need to be developed. In particular, Bayesian pair plots show strong correlation between seeding parameters and R_0, which can potentially lead again to a bimodal posterior where initial growth dynamics can be explained through R_0 or via initial seed infections.

8 Confounding and Causality: Estimating the Effect of Interventions

Section 6 showed that changes in the reproduction number over time can be explained by parameterising it in terms of covariates, such as NPIs or mobility. Clarifying the effects of interventions on disease transmission is important to guide policy and because NPIs have large economic and human consequences. In practice, effect sizes may not be identifiable for various reasons. NPIs often occurred in quick succession or simultaneously, leading to collinearity. They could be confounded with unobserved behavioural changes. Finally, the random distribution for the time between an infection and its recording as a case or death adds a large amount of variance to the observed data—one should not expect to see sharp or immediate changes, especially in reported deaths.

Flaxman et al. (2020b) estimated the effectiveness of NPIs across 11 European countries, and used partial pooling of effect sizes to address the identification problem. At that time, little data existed other than information on deaths and the timing of interventions. NPIs, which were coded as a binary set of mandatory government measures (e.g. school closures, ban on public events, lockdown), could not fully explain the patterns seen in some countries (e.g. Sweden), and especially at the subnational level. Mobility data became available in April and was used to model the epidemic in Italy, Brazil and the USA (Vollmer et al., 2020; Mellan et al., 2020; Unwin et al., 2020b). Such data is useful as it may help account for behavioral changes that confound the effects of NPIs. However since mobility affects transmission, is linked to the introduction of NPIs and potentially also to voluntary behavioural measures, we expect it to be a confounder. Sharma et al. (2021) further allow for a residual stochastic process (a random walk) to be included alongside the fixed NPI effects and perform estimation at a subnational level using a randomised study design.

Section 8.2 extends the model in Flaxman et al. (2020b) to further investigate this issue of confounding, and models both NPIs and mobility jointly. This is in keeping with standard practice in regression/ANOVA: expanding a model to take into account more explanatory variables. Nonetheless, NPIs may partially affect transmission *via a pathway through mobility*. A joint model of mobility and NPIs does not account for this. Therefore, in Section 8.3 we take a first and basic step in assessing causal considerations through a simple mediation analysis. We begin however by exploring the relationship between interventions and mobility.

8.1 Interventions and Mobility

Here we study the first epidemic wave in 2020, the same period considered by Flaxman et al. (2020b). We consider the simple case of regressing average mobility on the NPIs defined in Flaxman et al. (2020b), asking whether the changes in mobility can be explained by the timing of NPIs? Regressing average mobility on NPIs in a Bayesian linear model (no intercept or partial pooling) we find a correlation (Pearson's) of over 85% with a mean absolute error of 0.1%. Given mobility generally ranges from -1 to 1, this is a good overall fit. Figure 1a shows that visually, these fits correspond well with changes in average mobility. One could conjecture that mobility and NPIs are lagged, but lagging NPI dates either forwards or backwards in time does not result in a better fitting model, see Figure 1(b). Indeed, Figure 1(b) supports the hypothesis that the timing of NPIs and changes in mobility are coupled. The coefficient sizes from this regression are consistent with the (Flaxman et al., 2020b) finding that the NPI with the largest effect size is lockdown, see Figure 1(c). We note that the definition of lockdown encompasses a number of specific interventions including closing work places and stores, banning gatherings of various sizes, stay-at-home orders, and more. A more thorough analysis would included fine-grained intervention definitions as in Sharma et al. (2021). While this regression analysis does not model transmission or the trajectory of the epidemic, it provides evidence for the consistency of mobility and NPIs, and the potential for confounding. For regularisation we used a hierarchical shrinkage prior (Piironen and Vehtari, 2017) that performs both shrinkage and variable selection simultaneously.

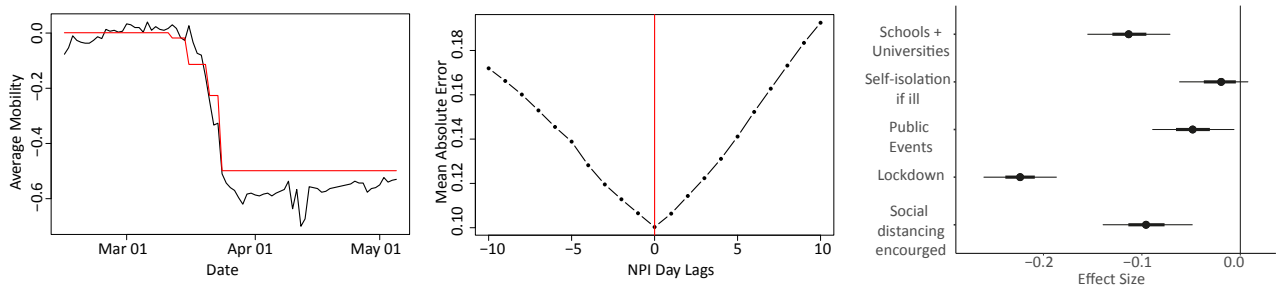

(a) (b) (c)

Figure 1: Simple regression of mobility against NPIs. (a) Regression prediction for the United Kingdom with mobility in black and the fit for NPIs in red. (b) The effect on mean absolute error from lagging the NPIs 10 days forward and backwards. (c) The coefficient effect sizes from the regression, with NPIs on the y-axis and regression effect sizes on the x-axis.

8.2 Controlling for Mobility

Section 8.1 found a correlation between interventions and mobility, demonstrating that mobility is a potential confounder. Here we control for this by jointly modeling NPIs and mobility. This is done using the same 11 European countries, sets of NPIs and death data as used in Flaxman et al. (2020b).

A two-stage approach (Haug et al., 2020) is used, whereby latent R_t is first estimated through a non-parametric daily random walk, i.e. independent of NPIs. We first nonparametrically estimate R_t to account for the various lags and biases in the observed data. In the second stage, R_t is regressed on NPIs and mobility. The random walk can in theory select any arbitrary function for R_t that best describes the data without any prior information on which interventions happened, when, or how well they worked. Given these estimates of R_t for all 11 European countries, we run a simple partial pooling model to see if interventions and/or mobility can reproduce the trends in R_t. The model used is a linear regression with country level intercepts (to account for variation in R_0), and both joint and country specific effect sizes for interventions/mobility. As with the earlier analysis we use a hierarchical shrinkage prior on the coefficients (Piironen and Vehtari, 2017).

Three variations of the model are considered: NPIs only, mobility only, and NPIs and mobility together. MCMC convergence diagnostics in all cases did not indicate fitting problems. We found the best fitting models (via PSIS-LOO (Vehtari et al., 2017)) to be NPIs alone or NPIs and mobility together. Relative to the NPIs and mobility together model the expected log posterior difference ($\pm$ standard error) in WAIC of the model with only NPIs is -5.2 ± 4 (not significant), and -565.6 ± 49.2 (significant) with only mobility. Therefore, in fits to the estimated R_t, the model with mobility alone is substantially worse than the models with NPIs. Controlling for mobility does not appear to significantly change the relative ranking of the estimated NPI effects. As in Flaxman et al. (2020b), the largest effect size is attributed to lockdown, as seen in Figure 2. This is true with and without the inclusion of the mobility variable.

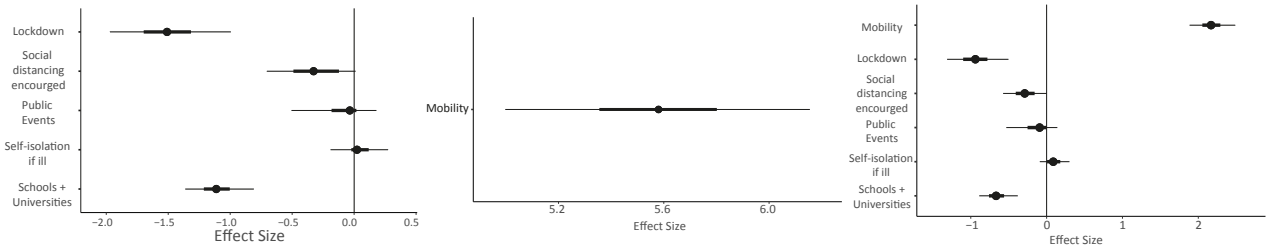

(a) (b) (c)

Figure 2: Regression of NPIs and/or mobility against nonparametric R_t. (a) NPIs only. (b) mobility only. (c) NPIs and mobility together. Mobility only was not *significantly* preferred by WAIC. Y axis are covariates and X axis the regression effect sizes

An advantage of the two-stage approach is that it is scaleable to a large number of regions (e.g. Laydon et al. (2021); Nouvellet et al. (2021)). R_t can be estimated in each region in parallel using separate models. Partial pooling can still

be leveraged to estimate effects in the second stage. Once R_t has been estimated, any number of interesting statistical analyses can be conducted. Nonetheless the estimated R_t is not entirely non-parametric; it is clearly influenced by the choice of using a first order random walk in the first stage. This analysis could be extended by considering a range of alternative priors for R_t, such as Gaussian processes. More importantly, however, this approach has not considered causal relationships between NPIs and mobility. This is the focus of the next example.

8.3 Causal Mediation

We would expect that part of the effect of NPIs on transmission occurs indirectly through its effect on mobility. If we *a priori* hypothesise that changes in mobility are both an effect of NPIs and a cause of reductions in transmission, causal mediation analysis (Pearl (2009)) provides a simple means to disentangle the total effect of a variable into a direct and indirect effect. The indirect effect occurs via some mediator, which in this case is hypothesised to be mobility.

Here we consider lockdown on its own, because performing causal mediation with all NPIs is challenging and lockdown is consistently the NPI with the largest effect size as shown above and in Flaxman et al. (2020b). As previously mentioned, the definition of lockdown represents an aggregate of policies, varying between countries. This analysis is therefore simply illustrative rather than being fully exhaustive. Briefly, to perform causal mediation we consider two transmission models

$$R_t^{(m)} = \tilde{R}_m^1 \exp\left(\left(\beta_1^1 + \beta_{1,m}^1\right) L_{t,m} + \epsilon_{t,m}^1\right), \tag{12}$$

$$R_t^{(m)} = \tilde{R}_m^2 \exp\left(\left(\beta_1^2 + \beta_{1,m}^2\right) L_{t,m} + \left(\beta_2^2 + \beta_{2,m}^2\right) M_{t,m} + \epsilon_{t,m}^2\right), \tag{13}$$

where $L_{t,m}$ is a binary indicator for lockdown and $M_{t,m}$ is mobility in country m respectively. $\tilde{R}_m^i$ and $\epsilon_{t,m}^i$ are country specific parameters modeling baseline transmission and a weekly random walk respectively. All other aspects of both models are the same as in Flaxman et al. (2020b). Model (12) includes effects for lockdown, while (13) additionally considers mobility. β_1^1 is the total effect for lockdown, while β_1^2 is the partial effect when controlling for mobility. The mediated effect is therefore $\beta_1^1 - \beta_1^2$. This quantifies the effect of lockdown *via the pathway through mobility*. We find this mediated effect reduces R_t by 18.3% with a 95% credible interval of $[12.2\%, 44.4\%]$. Individual coefficients are shown in Figure 3.

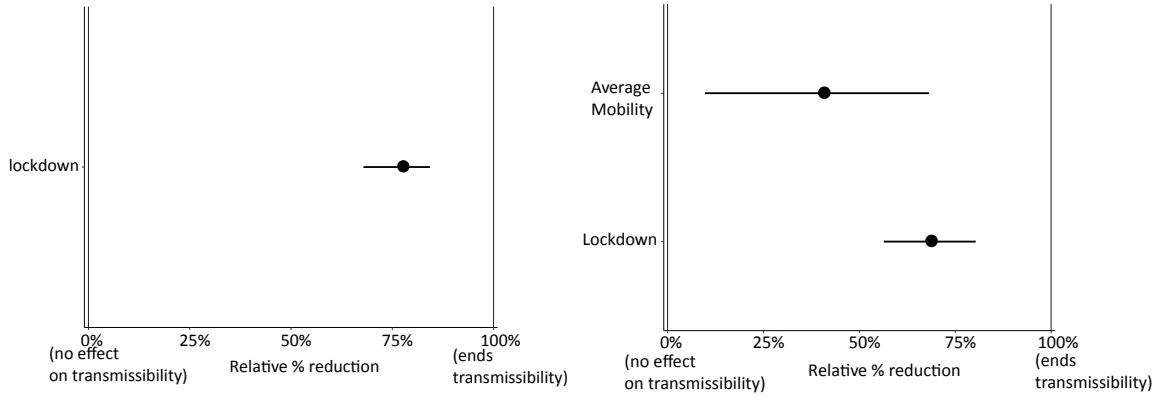

Figure 3: Mediation analysis. (a) The effect of lockdown from Model (12) (b) The effect of lockdown and mobility from Model (13)

These mediation results suggest a causal link between a lockdown policy and mobility that eventually leads to reduced transmission rates. They also suggest that the mediated effect is far less than the total effect of lockdown, suggesting lockdown will have other causal pathways. Of course, mobility is also mediated through other pathways, and a more extensive causal analysis is beyond the scope of this article. The exclusion of other NPIs may introduce omitted variable bias.

8.4 Simulation

We perform a simple simulation analysis to demonstrate the utility of our modelling framework. In these simulations we create artificial data with properties resembling the real data (e.g. intervention timings) but with hypothetical NPI effect sizes. We fitted our modelling framework to this simulated data to see if we could recover the hypothetical NPI effect sizes. We perform three simulation experiments where, to ensure consistency, the NPI dates are selected as those that actually happened in reality (Flaxman et al., 2020b). In selecting NPI timings from real data, we could create a more plausible representation of reality with the same orderings and collinearity. For all three simulations, R_0 and the initial epidemic seeding were set as those previously used (Flaxman et al., 2020b). The model used for estimating the effect sizes for all three simulations is the partial pooling model as described in (Flaxman et al., 2020c). The three simulation scenarios were as follows: (a) All 5 NPIs effect sizes were set to 20% - corresponding to the case where all interventions work equally well. (b) All 5 effect sizes were set to 18%, except there was an *unobserved* NPI which had an effect size of 10% and was applied at a random time at least 7 days before the last NPI occurred. This scenario corresponds to the case where there is an NPI that we did not account for in our model but which has an effect on transmission. (c) A single NPI was highly effective with an effect size of 70% and the remaining 30% were uniformly distributed among the remaining 4 NPIs. This scenario corresponds to a single important NPI that has the main effect in reducing transmission.

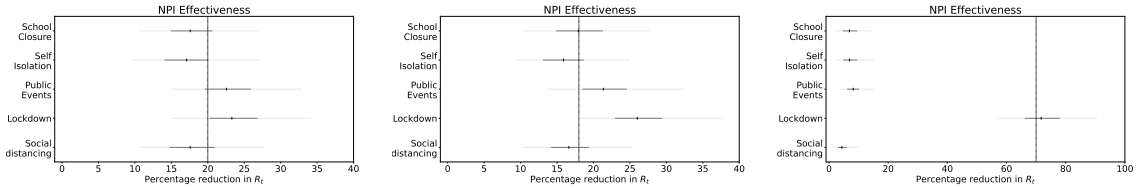

Figure 4: Simulation scenarios. (a) All NPIs have an effect size of 20%. (b) All NPIs have an effect size of 18% but there is an unobserved NPI with an effect size of 10%. (c) One NPI has an effect size of 70% and the remaining 30% is distributed to the 4 other NPIs

In Figure 4 we show that our model can recover the true effect sizes in all scenarios and, to a degree, motivates the appropriateness of our framework for modelling the effect of NPIs. However, our simulation experiments do not mean there are not significant shortcomings to our approach. A number of well known problems with statistical estimation apply to our model. Residual confounding from unobserved covariates (e.g. passive NPIs, behavioural changes) as in scenario (b) can be mitigated, but in general we recommend the inclusion of a stochastic process to model residual variation not captured by fixed effects (Sharma et al., 2021). Effect size estimates can be biased, even when residual variation is modelled. We recommend using debiasing techniques popular in statistics and machine learning (van der Laan and Rubin, 2006; Chernozhukov et al., 2017). Generalisation to different countries and times is not expected to be guaranteed (Sharma et al., 2021) and care should be taken when applying effect sizes to other countries and different time periods. Our modelling framework will inevitably be sensitive to decisions about fundamental epidemiological parameters, and these need to be regularly updated based on the best current evidence.

In summary, these simulation results show our framework is empirically motivated, and is a useful approach to estimating the effect of interventions on infectious disease epidemics. However, our approach does not surmount the common statistical problems affecting regression models in general. Care must be taken with critical decisions regarding data and modelling. . We believe our approach serves as a basic framework from which further development is needed (Sharma et al., 2021).

8.5 Code and Software

All analysis was run using the programming language R and the software package Stan (Carpenter et al., 2017). Code for transmission modelling is available from previously published studies (Scott et al., 2020; Flaxman et al., 2020b) and an online repository at `https://github.com/ImperialCollegeLondon/semi_mechanistic_renewal_processes`.

9 Discussion

This article has discussed a class of Bayesian semi-mechanistic statistical models for epidemics such as Covid-19 which are able to capture key epidemiological mechanisms. The model has appeared in various forms for specific analyses during the Covid-19 crisis and, at the time of writing, continues to be used to inform public policy. By presenting it in a general form and discussing key modelling difficulties we hope to stimulate discussion around it. As is constantly the predominant factor in empirical statistical approaches, the model is limited by data quality and availability. In this analysis we use a coarse definition of interventions that are likely to miss important details driving transmission. A key recommendation of future pandemic preparedness is to establish data pipelines that can quickly facilitate statistical modelling of the type outlined in this paper.

One key difficulty within the framework is dealing with confounded variables, particularly those used to explain changes in transmission during the early stages of an epidemic. The analyses in Section 8 make a first step in dealing with these. A number of model enhancements have not been included here and are an important area for further research. These include explicitly accounting for importations, allowing for uncertainty in the generation and infection-to-observation distributions, more expressive causal model, and the inclusion of residual effects using stochastic processes. The presented model can readily be fit using probabilistic programming languages such as Stan (Stan Development Team, 2018), though we note that the adaptive Hamiltonian Monte Carlo algorithm can at times face convergence problems when latent infections are modeled directly, or when multiple regions are jointly modeled. We conjecture that convergence may be improved by carefully choosing initial parameters for the sampler. Future research could explore whether alternative samplers can be developed to fit these models more pragmatically.

References

Badr, H. S., Du, H., Marshall, M., Dong, E., Squire, M. M. and Gardner, L. M. (2020) Association between mobility patterns and COVID-19 transmission in the USA: a mathematical modelling study. *Lancet Infect. Dis.*, **20**, 1247–1254.

Bellman, R. and Harris, T. (1948) On the theory of age-dependent stochastic branching processes. *Proceedings of the National Academy of Sciences*, **34**, 601–604.

Bettencourt, L. M. and Ribeiro, R. M. (2008) Real time Bayesian estimation of the epidemic potential of emerging infectious diseases. *PLoS ONE*, **3**, e2185.

Carias, C., O'Hagan, J. J., Gambhir, M., Kahn, E. B., Swerdlow, D. L. and Meltzer, M. I. (2019) Forecasting the 2014 West African Ebola Outbreak. *Epidemiologic Reviews*, **41**, 34–50.

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P. and Riddell, A. (2017) Stan: A probabilistic programming language. *J. Stat. Softw.*, **76**, 1–32.

Cauchemez, S., Valleron, A. J., Boëlle, P. Y., Flahault, A. and Ferguson, N. M. (2008) Estimating the impact of school closure on influenza transmission from Sentinel data. *Nature*, **452**, 750–754.

Champredon, D., Dushoff, J. and Earn, D. J. (2018) Equivalence of the Erlang-distributed SEIR epidemic model and the renewal equation. *SIAM Journal on Applied Mathematics*, **78**, 3258–3278.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C. and Newey, W. (2017) Double/Debiased/Neyman machine learning of treatment effects. *Am. Econ. Rev.*, **107**, 261–265.

Cori, A., Ferguson, N. M., Fraser, C. and Cauchemez, S. (2013) A new framework and software to estimate time-varying reproduction numbers during epidemics. *American Journal of Epidemiology*, **178**, 1505–1512.

Cowling, B. J., Ali, S. T., Ng, T. W., Tsang, T. K., Li, J. C., Fong, M. W., Liao, Q., Kwan, M. Y., Lee, S. L., Chiu, S. S., Wu, J. T., Wu, P. and Leung, G. M. (2020) Impact assessment of non-pharmaceutical interventions against coronavirus disease 2019 and influenza in Hong Kong: an observational study. *The Lancet Public Health*.

Crump, K. and Mode, C. J. (1969) A general age-dependent branching process. II. *J. Math. Anal. Appl.*, **25**, 8–17.

Ferguson, N. M., Donnelly, C. A. and Anderson, R. M. (2001) Transmission intensity and impact of control policies on the foot and mouth epidemic in Great Britain. *Nature*.

Flaxman, S., Mishra, S., Gandy, A., Unwin, H., Coupland, H., Mellan, T., Zhu, H., Berah, T., Eaton, J., Perez Guzman, P., Schmit, N., Cilloni, L., Ainslie, K., Baguelin, M., Blake, I., Boonyasiri, A., Boyd, O., Cattarino, L., Ciavarella, C., Cooper, L., Cucunuba Perez, Z., Cuomo-Dannenburg, G., Dighe, A., Djaafara, A., Dorigatti, I., Van Elsland, S., Fitzjohn, R., Fu, H., Gaythorpe, K., Geidelberg, L., Grassly, N., Green, W., Hallett, T., Hamlet, A., Hinsley, W., Jeffrey, B., Jorgensen, D., Knock, E., Laydon, D., Nedjati Gilani, G., Nouvellet, P., Parag, K., Siveroni, I., Thompson, H., Verity, R., Volz, E., Walters, C., Wang, H., Wang, Y., Watson, O., Winskill, P., Xi, X., Whittaker, C., Walker, P., Ghani, A., Donnelly, C., Riley, S., Okell, L., Vollmer, M., Ferguson, N. and Bhatt, S. (2020a) Report 13: Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in 11 european countries. *Tech. rep.*

Flaxman, S., Mishra, S., Gandy, A., Unwin, H. J. T., Mellan, T. A., Coupland, H., Whittaker, C., Zhu, H., Berah, T., Eaton, J. W., Monod, M., Perez-Guzman, P. N., Schmit, N., Cilloni, L., Ainslie, K. E. C., Baguelin, M., Boonyasiri, A., Boyd, O., Cattarino, L., Cooper, L. V., Cucunubá, Z., Cuomo-Dannenburg, G., Dighe, A., Djaafara, B., Dorigatti, I., van Elsland, S. L., FitzJohn, R. G., Gaythorpe, K. A. M., Geidelberg, L., Grassly, N. C., Green, W. D., Hallett, T., Hamlet, A., Hinsley, W., Jeffrey, B., Knock, E., Laydon, D. J., Nedjati-Gilani, G., Nouvellet, P., Parag, K. V., Siveroni, I., Thompson, H. A., Verity, R., Volz, E., Walters, C. E., Wang, H., Wang, Y., Watson, O. J., Winskill, P., Xi, X., Walker, P. G. T., Ghani, A. C., Donnelly, C. A., Riley, S. M., Vollmer, M. A. C., Ferguson, N. M., Okell, L. C., Bhatt, S. and Team, I. C. C.-. R. (2020b) Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature*. URL: `https://doi.org/10.1038/s41586-020-2405-7`.

Flaxman, S., Mishra, S., Scott, J., Ferguson, N., Gandy, A. and Bhatt, S. (2020c) Reply to: The effect of interventions on COVID-19. *Nature*, **588**, E29–E32.

Fraser, C. (2007) Estimating individual and household reproduction numbers in an emerging epidemic. *PLoS ONE*.

Fraser, C., Donnelly, C. A., Cauchemez, S., Hanage, W. P., Van Kerkhove, M. D., Hollingsworth, T. D., Griffin, J., Baggaley, R. F., Jenkins, H. E., Lyons, E. J., Jombart, T., Hinsley, W. R., Grassly, N. C., Balloux, F., Ghani, A. C., Ferguson, N. M., Rambaut, A., Pybus, O. G., Lopez-Gatell, H., Alpuche-Aranda, C. M., Chapela, I. B., Zavala, E. P., Ma. Espejo Guevara, D., Checchi, F., Garcia, E., Hugonnet, S. and Roth, C. (2009) Pandemic potential of a strain of influenza A (H1N1): Early findings. *Science*.

Gelman, A. and Hill, J. (2006) *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press.

Haug, N., Geyrhofer, L., Londei, A., Dervic, E., Desvars-Larrive, A., Loreto, V., Pinior, B., Thurner, S. and Klimek, P. (2020) Ranking the effectiveness of worldwide COVID-19 government interventions. *MedRxiv*.

of HKSAR, S. E. C. (2003) Chronology of the SARS epidemic in Hong Kong. *Tech. rep.*, Hong Kong.

Hox, J. J., Moerbeek, M. and de Schoot, R. (2010) *Multilevel analysis: Techniques and applications*. Routledge.

Islam, N., Sharp, S. J., Chowell, G., Shabnam, S., Kawachi, I., Lacey, B., Massaro, J. M., D'Agostino, Sr, R. B. and White, M. (2020) Physical distancing interventions and incidence of coronavirus disease 2019: natural experiment in 149 countries. *BMJ*, **370**, m2743.

Jagers, P. (1969) A general stochastic model for population development. *Scand. Actuar. J.*, **1969**, 84–103.

Kelly, H. A., Mercer, G. N., Fielding, J. E., Dowse, G. K., Glass, K., Carcione, D., Grant, K. A., Effler, P. V. and Lester, R. A. (2010) Pandemic (H1N1) 2009 influenza community transmission was established in one Australian state when the virus was first identified in North America. *PLoS ONE*.

Kermack, William Ogilvy and McKendrick, A. G. (1927) A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*.

Kimmel, M. (1983) The point-process approach to age- and time-dependent branching processes.

Kreft, I. and de Leeuw, J. (2011) *Introducing Multilevel Modeling*.

van der Laan, M. J. and Rubin, D. (2006) Targeted maximum likelihood learning. *Int. J. Biostat.*, **2**.

Laydon, D. J., Mishra, S., Hinsley, W. R., Samartsidis, P., Flaxman, S., Gandy, A., Ferguson, N. M. and Bhatt, S. (2021) Modelling the impact of the tier system on SARS-CoV-2 transmission in the UK between the first and second national lockdowns. *BMJ Open*, **11**, e050346.

Li, R., Pei, S., Chen, B., Song, Y., Zhang, T., Yang, W. and Shaman, J. (2020) Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV-2). *Science*.

Lloyd-Smith, J. O., Schreiber, S. J., Kopp, P. E. and Getz, W. M. (2005) Superspreading and the effect of individual variation on disease emergence. *Nature*, **438**, 355–359.

May, R. M. (1976) Simple mathematical models with very complicated dynamics. *Nature*.

Mellan, T. A., Hoeltgebaum, H. H., Mishra, S., Whittaker, C., Schnekenberg, R. P., Gandy, A., Unwin, H. J. T., Vollmer, M. A. C., Coupland, H., Hawryluk, I., Faria, N. R., Vesga, J., Zhu, H., Hutchinson, M., Ratmann, O., Monod, M., Ainslie, K., Baguelin, M., Bhatia, S., Boonyasiri, A., Brazeau, N., Charles, G., Cooper, L. V., Cucunuba, Z., Cuomo-Dannenburg, G., Dighe, A., Djaafara, B., Eaton, J., van Elsland, S. L., FitzJohn, R., Fraser, K., Gaythorpe, K., Green, W., Hayes, S., Imai, N., Jeffrey, B., Knock, E., Laydon, D., Lees, J., Mangal, T., Mousa, A., Nedjati-Gilani, G., Nouvellet, P., Olivera, D., Parag, K. V., Pickles, M., Thompson, H. A., Verity, R., Walters, C., Wang, H., Wang, Y., Watson, O. J., Whittles, L., Xi, X., Okell, L., Dorigatti, I., Walker, P., Ghani, A., Riley, S., Ferguson, N. M., Donnelly, C. A., Flaxman, S. and Bhatt, S. (2020) Report 21 - estimating COVID-19 cases and reproduction number in brazil.

Miller, A. C., Foti, N. J., Lewnard, J. A., Jewell, N. P., Guestrin, C. and Fox, E. B. (2020) Mobility trends provide a leading indicator of changes in SARS-CoV-2 transmission. *medRxiv*. URL: `https://www.medrxiv.org/content/early/2020/05/11/2020.05.07.20094441`.

Mishra, S., Scott, J., Zhu, H., Ferguson, N. M., Bhatt, S., Flaxman, S. and Gandy, A. (2020) A covid-19 model for local authorities of the united kingdom. *medRxiv*. URL: `https://www.medrxiv.org/content/early/2020/11/27/2020.11.24.20236661`.

Nouvellet, P., Bhatia, S., Cori, A., Ainslie, K. E. C., Baguelin, M., Bhatt, S., Boonyasiri, A., Brazeau, N. F., Cattarino, L., Cooper, L. V., Coupland, H., Cucunuba, Z. M., Cuomo-Dannenburg, G., Dighe, A., Djaafara, B. A., Dorigatti, I., Eales, O. D., van Elsland, S. L., Nascimento, F. F., FitzJohn, R. G., Gaythorpe, K. A. M., Geidelberg, L., Green, W. D., Hamlet, A., Hauck, K., Hinsley, W., Imai, N., Jeffrey, B., Knock, E., Laydon, D. J., Lees, J. A., Mangal, T., Mellan, T. A., Nedjati-Gilani, G., Parag, K. V., Pons-Salort, M., Ragonnet-Cronin, M., Riley, S., Unwin, H. J. T., Verity, R., Vollmer, M. A. C., Volz, E., Walker, P. G. T., Walters, C. E., Wang, H., Watson, O. J., Whittaker, C., Whittles, L. K., Xi, X., Ferguson, N. M. and Donnelly, C. A. (2021) Reduction in mobility and COVID-19 transmission. *Nat. Commun.*, **12**, 1090.

Nouvellet, P., Cori, A., Garske, T., Blake, I. M., Dorigatti, I., Hinsley, W., Jombart, T., Mills, H. L., Nedjati-Gilani, G., Van Kerkhove, M. D., Fraser, C., Donnelly, C. A., Ferguson, N. M. and Riley, S. (2018) A simple approach to measure transmissibility and forecast incidence. *Epidemics*.

NYS Press Office (2020) Amid ongoing covid-19 pandemic, governor cuomo announces state is bringing in international experts to help advise the state's reopening plan. URL: `https://www.governor.ny.gov/news/amid-ongoing-covid-19-pandemic-governor-cuomo-announces-state-bringing-international-exp`

Olney, A. M., Smith, J., Sen, S., Thomas, F. and Unwin, H. J. T. (2020) Estimating the effect of social distancing interventions on covid-19 in the united states. *medRxiv*.

Pakkanen, M. S., Miscouridou, X., Berah, T., Mishra, S., Mellan, T. A. and Bhatt, S. (2022) Unifying incidence and prevalence under a time-varying general branching process. ArXiv:2107.05579 [q-bio.PE].

Pearl, J. (2009) Causal inference in statistics: An overview. *Stat. Surv.*, **3**, 96–146.

Piironen, J. and Vehtari, A. (2017) Sparsity information and regularization in the horseshoe and other shrinkage priors. *Electron. J. Stat.*, **11**, 5018–5051.

Riley, S., Fraser, C., Donnelly, C. A., Ghani, A. C., Abu-Raddad, L. J., Hedley, A. J., Leung, G. M., Ho, L. M., Lam, T. H., Thach, T. Q., Chau, P., Chan, K. P., Lo, S. V., Leung, P. Y., Tsang, T., Ho, W., Lee, K. H., Lau, E. M., Ferguson, N. M. and Anderson, R. M. (2003) Transmission dynamics of the etiological agent of SARS in Hong Kong: Impact of public health interventions. *Science*.

Rizoiu, M.-A., Xie, L., Sanner, S., Cebrian, M., Yu, H. and Van Hentenryck, P. (2017) Expecting to be hip: Hawkes intensity processes for social media popularity. In *Proceedings of the 26th International Conference on World Wide Web*, WWW '17, 735–744. Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee. URL: `https://doi.org/10.1145/3038912.3052650`.

Scott, J. A., Gandy, A., Mishra, S., Unwin, J., Flaxman, S. and Bhatt, S. (2020) epidemia: Modeling of Epidemics using Hierarchical Bayesian Models. URL: `https://imperialcollegelondon.github.io/epidemia/`.

Sharma, M., Mindermann, S., Rogers-Smith, C., Leech, G., Snodin, B., Ahuja, J., Sandbrink, J. B., Monrad, J. T., Altman, G., Dhaliwal, G., Finnveden, L., Norman, A. J., Oehm, S. B., Sandkühler, J. F., Aitchison, L., Gavenčiak, T., Mellan, T., Kulveit, J., Chindelevitch, L., Flaxman, S., Gal, Y., Mishra, S., Bhatt, S. and Brauner, J. M. (2021) Understanding the effectiveness of government interventions against the resurgence of COVID-19 in europe. *Nat. Commun.*, **12**, 5820.

Stan Development Team (2018) The {Stan} Core Library. URL: `http://mc-stan.org/`.

The Scottish Government (2020) Coronavirus (covid-19): modelling the epidemic. URL: `https://www.gov.scot/collections/coronavirus-covid-19-modelling-the-epidemic/`.

Thompson, W. A., Taylor, H. M. and Karlin, S. (1984) *An Introduction to Stochastic Modeling*. Academic Press.

Tokuda, T., Goodrich, B., Van Mechelen, I., Gelman, A. and Tuerlinckx, F. (2011) Visualizing distributions of covariance matrices. *Columbia Univ., New York, USA, Tech. Rep*, 18–18.

Unwin, H., Mishra, S., Bradley, V. C., Gandy, A., Vollmer, M. and others (2020a) Report 23: State-level tracking of COVID-19 in the united states.

Unwin, H. J. T., Mishra, S., Bradley, V. C., Gandy, A., Mellan, T. A., Coupland, H., Ish-Horowicz, J., Vollmer, M. A. C., Whittaker, C., Filippi, S. L., Xi, X., Monod, M., Ratmann, O., Hutchinson, M., Harrison Zhu, F. V., Hawryluk, I., Milton, P., Ainslie, K. E. C., Baguelin, M., Boonyasiri, A., Brazeau, N. F., Cattarino, L., Cucunuba, Z., Cuomo-Dannenburg, G., Dorigatti, I., Eales, O. D., Eaton, J. W., van Elsland, S. L., FitzJohn, R. G., Gaythorpe, K. A. M., Green, W., Hinsley, W., Jeffrey, B., Knock, E., Laydon, D. J., Lees, J., Nedjati-Gilani, G., Nouvellet, P., Okell, L., Parag, K. V., Siveroni, I., Thompson, H. A., Walker, P., Walters, C. E., Watson, O. J., Whittles, L. K., Ghani, A. C., Ferguson, N. M., Riley, S., Donnelly, C. A., Bhatt, S. and Flaxman, S. (2020b) State-level tracking of COVID-19 in the united states. *Nat. Commun.*

Vehtari, A., Gelman, A. and Gabry, J. (2017) Practical bayesian model evaluation using leave-one-out cross-validation and WAIC. *Stat. Comput.*, **27**, 1413–1432.

Vollmer, M. A. C., Mishra, S., Unwin, H. J. T., Gandy, A., Mellan, T. A., Bradley, V., Zhu, H., Coupland, H., Hawryluk, I., Hutchinson, M., Ratmann, O., Monod, M., Walker, P., Whittaker, C., Cattarino, L., Ciavarella, C., Cilloni, L., Ainslie, K., Baguelin, M., Bhatia, S., Boonyasiri, A., Brazeau, N., Charles, G., Cooper, L. V., Cucunuba, Z., Cuomo-Dannenburg, G., Dighe, A., Djaafara, B., Eaton, J., van Elsland, S. L., FitzJohn, R., Fraser, K., Gaythorpe, K., Green, W., Hayes, S., Imai, N., Jeffrey, B., Knock, E., Laydon, D., Lees, J., Mangal, T., Mousa, A., Nedjati-Gilani, G., Nouvellet, P., Olivera, D., Parag, K. V., Pickles, M., Thompson, H. A., Verity, R., Walters, C., Wang, H., Wang, Y., Watson, O. J., Whittles, L., Xi, X., Ghani, A. A. M., Riley, S., Okell, L. C., Donnelly, C. A., Ferguson, N. M., Dorigatti, I., Flaxman, S. and Bhatt, S. (2020) A sub-national analysis of the rate of transmission of COVID-19 in italy. *medRxiv*.

Wallinga, J. and Teunis, P. (2004) Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures. *American Journal of Epidemiology*, **160**, 509–516.

WHO (2003) SARS: chronology of a serial killer. *Tech. rep.*

10 Appendix

10.1 Offspring Dispersion

Define the offspring distribution of any given infection to be the distribution of the random number of offspring attributable to that infection. We show that assuming the variance of these distributions are a constant proportion of the mean implies, under suitable independence assumptions, the same result for new infections I_t for all time points.

Assume some ordering over infections at each period, and let $O_t^{(i)}$ denote the number of offspring of the i^{th} infection at time t. This can be decomposed as

$$O_t^{(i)} = \sum_{s=t+1}^{\infty} O_{ts}^{(i)}, \tag{14}$$

where $O_{ts}^{(i)}$ are the number of offspring of i birthed at time s. The branching process behind Equation (5) implies that $O_{ts}^{(i)}$ has mean $R_s g_{s-t}$. Assume that $\{O_{ts}^{(i)} : s \geq t\}$ are mutually independent and have variance which is a fixed proportion d of the mean. By Equation (14), this implies the same variance relationship for $O_t^{(i)}$. In particular, if $R_s = R_t$ for $s > t$ then $O_t^{(i)}$ has mean R_t and variance dR_t. New infections at time t can be expressed as

$$I_t = \sum_{s=1}^{t-1} \sum_{i=1}^{I_s} O_{st}^{(i)}. \tag{15}$$

Assume that all $O_{st}^{(i)}$ appearing in Equation (15) are mutually independent conditional on everything occurring up to time $t-1$, the result clearly follows by taking the variance of both sides of Equation 15 given R_t and $I_{v:t-1}$.

10.2 Population Adjustment

Here we motivate Equation (6), which is used to adjust transmission rates for the size of the infectable population. The most obvious starting point for such an adjustment would be to let

$$\mathrm{E}[I_t | R_t, I_{v:t-1}] = \left(\frac{S_0 - I_{t-1}}{S_0} \right) R_{u,t} L_t, \tag{16}$$

where R_{u_t} is defined as in Section 4.2. This is similar in form to a *discrete logistic growth model*. Such models are well known as examples of simple models that exhibit chaotic dynamics (May, 1976). In particular, it is possible that the expected value on the left hand side exceeds the remaining susceptible population. Intuitively, this issue occurs because multiple infections can occur simultaneously in the discrete model. We therefore propose solving this by using a population adjustment motivated by the solution to a continuous time model whose intensity is a simplification of Equation (3).

Suppose we observe $I_{v:t-1}$ and current transmission R_t. We evolve infections from time $t-1$ to t continuously, and hence avoid overshooting. Define a continuous time counting $\tilde{I}(s)$ process starting at time $t-1$ by the intensity

$$\tilde{\lambda}(s) = \left(\frac{S_0 - \tilde{I}(s)}{S_0} \right) R_{u,t} L_t, \tag{17}$$

for $s \geq t-1$, and with initial condition $\tilde{I}(t-1) = I_{t-1}$. Supplementary 11.1 shows that

$$\mathrm{E}[\tilde{I}(t)] = I_{t-1} + (S_0 - I_{t-1}) \left(1 - \exp\left(-\frac{R_{u,t} L_t}{S_0} \right) \right), \tag{18}$$

which is the motivation for Equation (6).

11 Online Supplement

11.1 Proof of Equation (18)

Without loss of generality, we prove the result for time $t = 1$. The argument remains the same for all $t > 1$.

From (Thompson et al., 1984, Lemma 5.5), we have

$$\mathrm{E}[\tilde{I}(s)] = \tilde{I}(0) + \int_0^s \mathrm{E}[\tilde{\lambda}(l)]dl \text{ for } s \geq 0.$$

The following lemma derives an expression for the the expected intensity on the right hand side.

Lemma 11.1. *The expected intensity takes the form*

$$E[\tilde{\lambda}(s)] = \tilde{\lambda}(0) \exp\left(-\frac{R_{u,1}L_1}{S_0}s\right),$$

for all $s \geq 0$.

Proof of Lemma 11.1. Fix $s \geq 0$, some small $\Delta > 0$ and let $h(s) := \mathrm{E}[\tilde{\lambda}(s)]$. We have from Equation (17) that

$$h(s+\Delta) = \left(\frac{S_0 - \mathrm{E}[\tilde{I}(s+\Delta)]}{S_0}\right) R_{u,1}L_1. \tag{19}$$

We can write

$$\mathrm{E}[\tilde{I}(s+\Delta)|\tilde{\lambda}(s)] = \mathrm{E}[\tilde{I}(s)|\tilde{\lambda}(s)] + \tilde{\lambda}(s)\Delta + \mathcal{O}(\Delta),$$

and taking expectations on both sides,

$$\mathrm{E}[\tilde{I}(s+\Delta)] = \mathrm{E}[\tilde{I}(s)] + h(s)\Delta + \mathcal{O}(\Delta).$$

Substituting this into (19) and rearranging gives

$$\begin{aligned} h(s+\Delta) &= \left(\frac{S_0 - \mathrm{E}[\tilde{I}(s+\Delta)]}{S_0}\right) R_{u,1}L_1 - \frac{R_{u,1}L_1}{S_0}\left(h(s)\Delta + \mathcal{O}(\Delta)\right), \\ &= h(s) - \frac{R_{u,1}L_1}{S_0}\left(h(s)\Delta + \mathcal{O}(\Delta)\right). \end{aligned}$$

Rearranging gives

$$\frac{h(s+\Delta) - h(s)}{\Delta} = -\frac{R_{u,1}L_1}{S_0}\left(h(s) + \frac{\mathcal{O}(\Delta)}{\Delta}\right).$$

Taking the limit as $\Delta \to 0$ and rearranging gives the differential equation

$$\frac{h'(s)}{h(s)} = -\frac{R_{u,1}L_1}{S_0}.$$

Integrating both sides gives

$$\log(h(s)) = -\frac{R_{u,1}L_1}{S_0}s + C.$$

Using that $h(0) = \tilde{\lambda}(0)$ gives the constant $C = \log(\tilde{\lambda}(0))$. Plugging in yields the required result. □

Hence,

$$\begin{aligned}
\mathrm{E}[\tilde{I}(s)] &= I_0 + \tilde{\lambda}(0) \int_0^s \exp\left(-\frac{R_{u,1}L_1}{S_0} l\right) dl \\
&= I_0 + \tilde{\lambda}(0) \frac{S_0}{R_{u,1}L_1} \left(1 - \exp\left(-\frac{R_{u,1}L_1}{S_0} s\right)\right) \\
&= I_0 + (S_0 - \tilde{I}(s)) \left(1 - \exp\left(-\frac{R_{u,1}L_1}{S_0} s\right)\right).
\end{aligned}$$

Letting $s = 1$ gives the required result.